

**REMARKS**

Claims 1-27 are pending in this application, all of which have been rejected. As a result of the June 16, 2004 Office Action, claims 1-3, 7, 9, 10, 22, 23, 25, and 26 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Lerman (U.S. Patent No. 6,378,036) in view of Pence (U.S. Patent No. 6,279,074). Claims 4-6, 8, and 24 stand rejected under section 103(a) as being unpatentable over Lerman in view of Pence and Tucker (U.S. Patent No. 6,178,519). Claims 11, 12, and 14-20 stand rejected under section 103(a) as being unpatentable over Konshak in view of Lerman. Following entry of the amendment, claims 28, 29, and 30 will have been added, and claim 25 will have been cancelled without prejudice or disclaimer of the subject matter thereof.

A telephonic interview was held on August 11, 2004 between the undersigned, Examiner Ehichioya, and Primary Examiner Kindred. Claims 1, 2, 3, 7, 21, 26, and 27 were discussed in view of the Lerman reference. No agreement was reached. However, the undersigned was advised at the interview that the Examiner would closely consider the features that were discussed at the interview. Thus, applicants have summarized certain features of the claims below, and request reconsideration based on these features. Moreover, applicants respectfully submit that the foregoing discussion of the interview, together with the discussion below of the features that were discussed in the interview, satisfies the obligation to summarize the interview under 37 C.F.R. § 1.133(b).

The queue-per-media feature

Claims 1, 21, and 26-30 each recite features that, while not identical in scope or language, effectively call for each media to have its own queue of requests. This feature is not taught in any of the art that has been cited. As discussed in the telephone interview, this feature enables all of the recall requests that will require the use of a particular media to be aggregated on a single queue, thereby preventing the switching back and forth between different media that would be required if requests to recall data items from different media were placed on the same queue. The portion of Lerman that has been cited as teaching this feature shows three queues – a steady-state subscriber queue (SSQ), a new subscriber queue (NSQ), and an other request queue (ORQ) – that contain requests from different media.

While Lerman arguably shows that these requests are channeled into a plurality of “SDS” queues (col. 4, ll. 34-41), each SDS queue corresponds to a particular *disk drive* (see col. 4, l. 36), not to a particular disk (or other type of media). (While Lerman refers, at times, to these queues as “disk queues,” it is clear from context that Lerman is using the phrase “disk queue” as a shorthand for what should properly be called “disk drive queue,” since Lerman appears to refer to fixed disk media rather than removable media, and thus does not take into account the distinction between a disk and the drive on which it is mounted.)

In short, Lerman teaches, at best, that each drive has its own queue, and not that each media has its own queue. Thus, for these reasons applicants request reconsideration of claims 1, 21, 26, and 27.

Additionally, while applicants believe that claims 1, 21, 26, and 27 define over Lerman (and the other art cited), applicants have added claims 28-30 to more particularly point out the invention. In particular, claims 28-30 call for the media to be separately mountable on a plurality of drives, where each media continues to correspond to a particular queue regardless of which drive the media is mounted on. Since Lerman teaches that the queue is associated with a particular drive rather than a particular disk – and since Lerman does not address what happens to a queue if a media is moved from one drive to another – claims 28-30 define over the art cited and should be allowed.

#### The monotonically-increasing offsets feature

Claims 2, 16, 17, 18, and 24 each call for queues that have one or more sets of requests that are each organized in monotonically increasing sequences *with respect to their offsets into a media*. As discussed at the interview, an “offset” into a media is a way of describing a given data item’s location, and the above-referenced claims organize data in increasing order of these offsets. In the interview, the undersigned pointed out that the Examiner has applied references that teach, at best, organizing data in a sequence with respect to time-based deadlines, and it was indicated that the Examiner may have overlooked the fact that the sequences that are recited in the claims increase with respect to offset rather than time. Applicants respectfully submit that there is no cited reference that teaches organizing requests on a queue in increasing order of offset in the manner claimed, and applicants request reconsideration of the rejection of the claims that recite this feature.

The concurrent retrieval feature

Claims 3 and 11 each call for concurrent retrieval of data from two different media. As pointed out in response to the previous office action, the Examiner has principally relied on Lerman, col. 1, ll. 39-42 as teaching this “concurrent” feature. While this cited passage mentions the word “concurrently,” the use of that word refers to concurrent video service, not concurrent retrieval of data from two different media. Applicants submit that the cited portion of Lerman does not teach the feature of concurrent retrieval for which it has been cited, and thus request that the rejection of claims 3 and 11 be withdrawn.

The timestamp feature

Claims 7 and 23 call for timestamping the queues as part of the queue-creation process, and then selecting a particular queue for activation based on the timestamp. As discussed at the interview, the timestamp effectively represents the age of the queue, and thus considering the timestamp in deciding which queue to activate can help to prevent (although not entirely prevent) old queues from being ignored when deciding which queue to activate. The Examiner has cited Pence as teaching this timestamping feature. However, the portion of Pence that has been cited teaches, at best, the timestamping of individual requests, not entire queues. Timestamping the queue – and then considering the timestamp in deciding which queue to activate – is useful in a multi-queue system such as the one claimed. Since Pence uses only one queue, timestamping that queue would be meaningless.

Thus, applicants requests reconsideration of the rejection of claims 7 and 23.

No new matter

New claims 28-30 do not add new matter. These claims are supported by the originally-filed specification at least at FIGS. 4A and 5, page 16, lines 22-24, and page 18, lines 15-16.

Conclusion

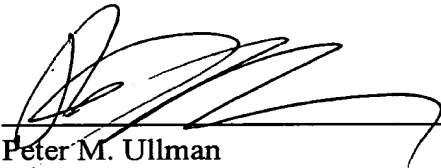
Claims 1, 2, 3, 7, 11, 16-18, 21, 23, 24, and 26-30 have been shown to be patentable for the reasons discussed above. Moreover, claims 4-6, 8-10, 12-15, 19, 20, and 22 are

**DOCKET NO.:** MSFT-0232/160298.01  
**Application No.:** 09/783,820  
**Office Action Dated:** June 16, 2004

**PATENT**

dependent – either directly or indirectly – on claims that have been shown to be patentable, and thus are patentable at least by reason of their dependency. Accordingly, applicants respectfully submit that this case is now in condition for allowance.

Date: September 7, 2004



---

Peter M. Ullman  
Registration No. 43,963

Woodcock Washburn LLP  
One Liberty Place - 46th Floor  
Philadelphia PA 19103  
Telephone: (215) 568-3100  
Facsimile: (215) 568-3439